White Paper

October 2022

A Deterministic Ethernet Solution based on Time Sensitive Networking (TSN)

Michael Hegarty (hegarty@ddc-web.com)
Product Line Manager, Data Bus Products

©2018 Data Device Corporation. All trademarks are the property of their respective owners.

Table of Contents

Table of Figures

1 Introduction

The recent introduction of the IEEE 802.1Q Time Sensitive Networking (TSN) real-time extensions to Ethernet presents new options for data networks within aerospace systems. Historically the widespread use of Ethernet in real-time systems has been hampered by Ethernet's "lossy" communication model and a general lack of determinism. TSN is an emerging technology for switched networks that provides high accuracy time synchronization, deterministic traffic scheduling and robust redundancy and path routing. These features combine to make it an attractive option for aerospace applications.

2 Overview of Time Sensitive Networking (TSN)

TSN, based on a series of IEEE 802.1Q standards, provides the high-level architecture for data queueing and scheduling methods, support of different data classes, redundancy, cyclic forwarding and queuing, frame preemption, time synchronization, redundancy, filtering and policing, and network configuration. These features combine to form the framework of a robust real-time network.

2.1 TSN: Introduction

For many years, Ethernet has been the widest deployed standard for home and office computing and is also widely used for industrial and military/aerospace applications. Ethernet's success has been attributable to it being an open standard with an extremely wide base of supplier support.

Over the years, the Ethernet supplier base has invested in evolutions for increased bandwidth, scalability, protocols, components, boards, and test equipment. There is also a large base of software written for Ethernet that includes drivers, stacks, APIs, middleware, support by many different operating systems and applications.

In 2005, the IEEE published the 802.1BA standard for Audio Video Bridging (AVB) Systems. This standard was targeted to consumer and professional audio and video applications.

In 2012, the IEEE formed the TSN task group, with the goal of expanding the use of AVB to applications beyond video and audio streaming. These real time embedded control and sensor systems support industrial automation, energy, automotive, aerospace, medical, and IoT (Internet of things) applications. The TSN group's charter was to define mechanisms for sending time sensitive information requiring deterministic behavior, guaranteed QoS (quality of service) and security over Ethernet networks.

Currently, TSN consists of a group of 12 standards that are built upon or linked to other Ethernet standards.

TSN does not specify a data rate and could be implemented over different types of networks operating at rates from 10 Mb/s to 100 Gb/s and higher. In addition to Ethernet, TSN can operate across other types of networks such as Wi-Fi.

Further, TSN can operate reliably over a variety of connection distances. Link distances can range from inches, such as within server racks; to feet, as in office buildings; to miles, for example on an automotive assembly line.

TSN adds a layer of determinism to Ethernet. Although there had been earlier implementations of deterministic Ethernet, many of these were proprietary and were targeted to specific application domains. The goal for TSN was to develop a set of standards providing a “tool box” that could be used in many different areas.

The major features of TSN are:

- Providing deterministic data delivery with guaranteed values of end-to-end latency.
- Providing support for “best effort” (e.g., like standard Ethernet), “rate-constrained” (e.g., ARINC-664/AFDX), and “scheduled” classes of data traffic that require determinism and minimal latency and latency jitter. Application developers can select the required traffic type(s) and mix for their system.
- Methods for “time aware” scheduling, including multiplexing data traffic from end points and through switches
- Defined redundancy methods for providing fault tolerance
- Methods for prioritizing different classes of data based on transmission time, thereby guaranteeing their forwarding and delivery at a defined point in time.
- Providing methods for synchronizing time across a network
- Providing methods for reserving slots with a pre-defined time cycle for transmitting periodic real-time data and reserving one or more paths through a network.

2.2 Types of Data Traffic

TSN supports three different classes of data traffic:

1. **Best Effort Traffic**. Best effort traffic is what's used on most Ethernet networks. That is, the network does not provide guaranteed delivery or a maximum delay time for delivering best effort messages to their destinations. With TSN, best effort traffic is assigned the lowest priority level. TSN scheduling software tools allow best effort traffic to be transmitted from an end point or a switch port output only during times when there's no rate constrained or scheduled traffic currently being transmitted. As a result, TSN does not provide any guarantees of delivery, latency or latency jitter for best effort traffic.

2. **Reserved Traffic**. Reserved traffic (sometimes referred to as rate-constrained traffic or credit-based shaped traffic) is similar to what's defined by the ARINC 664 Part 7 standard (also referred to as AFDX™) that's used as the backbone network for many commercial aircraft. For TSN, reserved traffic is assigned a higher priority than best effort traffic, but a lower priority than scheduled traffic. Most reserved messages are transmitted periodically by end points. Reserved traffic implements a credit-based shaper to constrain the data rate of traffic flow. The rate-constraints on the reserved traffic allows for the calculation of worst case latency and provides the framework for ensuring guaranteed delivery (assuming that the defined traffic rates do not exceed the bandwidth of the network). In addition, TSN includes a policing function in which switches monitor the volume of reserved traffic transmitted over a given stream in a given amount of time to verify that the packet rates do not exceed the defined maximum rate for a given traffic class. If a switch determines that an end point or upstream switch port is transmitting excessive amounts of data, it could make a determination to stop passing additional packets across the network. The assumption is that the network has a pre-established allocation of bandwidth among participating nodes. The policing function within the switch verifies compliance to the bandwidth allocation (implemented through the credit-based shapers), thus maintaining latency below the defined maximum and ensuring guaranteed delivery.

3. **Scheduled Traffic**. With TSN, scheduled traffic is assigned the highest priority level, higher than either best effort or rate constrained traffic. With scheduled traffic, the network provides a guarantee of delivery with a minimal latency. This is configured by the TSN scheduling tools, which guarantee a path and worst case end-to-end latency performance through the network.

2.3 TSN Scheduling

Figure 1 is an illustration of TSN's primary method for scheduling data traffic. This is defined by IEEE standard 802.1Qbv, Enhancements for Scheduled Traffic (which has been incorporated in the main 802.1Q standard in IEEE 802.1Q-2018). As shown, for transmitting from either an end point port or from the output port of a switch, the port maintains up to eight output queues of frames to be transmitted.

Each queue will be dedicated for either scheduled traffic, reserved (rate-constrained / credit-based) traffic or best effort traffic. Since best effort traffic is by definition the least urgent, all best effort traffic is typically stacked in the same queue and is given the lowest priority.

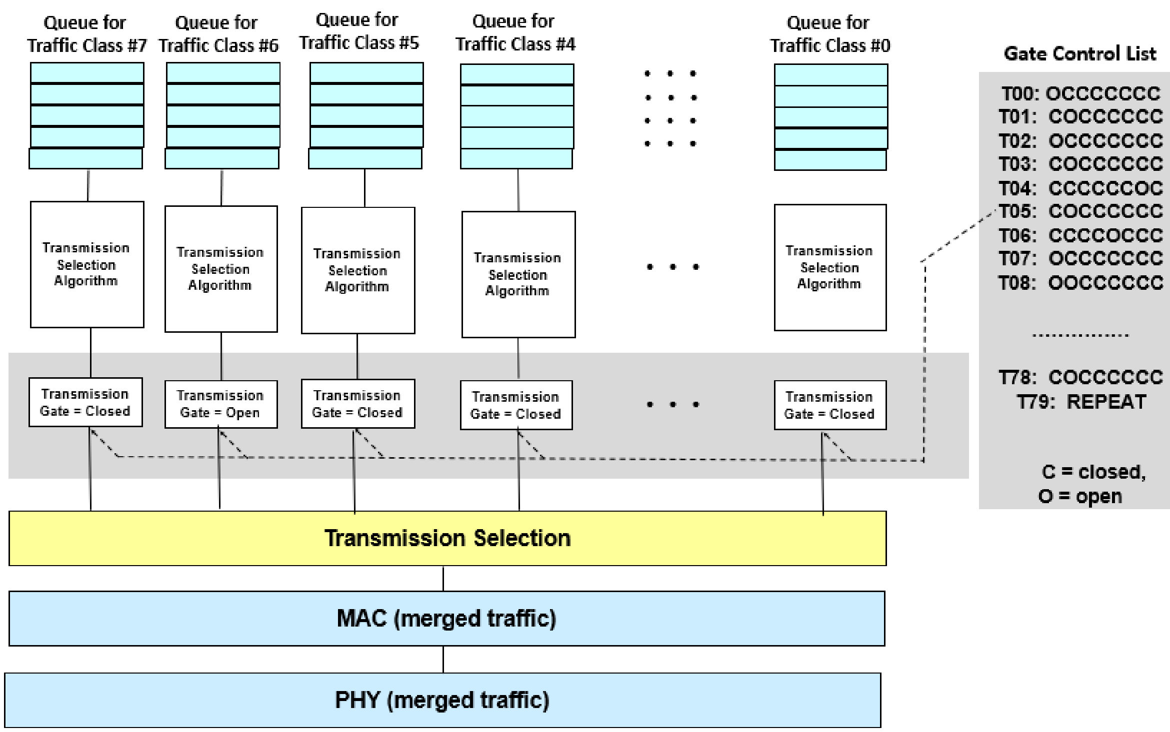

Figure 1. TSN Traffic Scheduler

Typically, the other 7 queues may be used for transmitting either scheduled or rate constrained traffic. The network scheduler will only allow best effort traffic to be transmitted on a non-interfering basis with either scheduled or rate constrained traffic.

For each queue, the end point or switch implements an algorithm for ordering the traffic to be transmitted within that particular queue. For a scenario where the number of input/output “stream pairs” is greater than eight, the transmission selection algorithm would need to stack frames for multiple streams on the same output queue (i.e. within the same traffic class).

The separate output queues are used for multiplexing between multiple streams of data traffic. For a given end point or switch output port, only one queue will be scheduled to transmit traffic at a given point in time. As shown in Figure 1, this queue’s transmission gate will be designated as “open”, while the transmission gate for the other seven queues will be designated as “closed”.

Each port’s logic will maintain a gate control list. The gate control list, which is created by the network scheduling software, lists a repetitive schedule that indicates when each of the individual queues is designated as the transmitting queue.

For TSN scheduled traffic, the scheduling software ensures non-blocking performance, with guaranteed end-to-end delivery and worst case latency time across a network. It does this by avoiding cases where multiple output ports from either the same switch or different switches on a network are scheduled to transmit to the same destination port (either an end point or switch port) at the same time.

2.4 Cyclic Queuing and Forwarding

For applications that can tolerate latency/latency jitter values on the order of 100 µS, a TSN technique known as Cyclic Queuing and Forwarding offers a highly reliable solution. This operation is defined by IEEE standard P802.1Qch.

Referring to Figure 2, cyclic queuing and forwarding specifies the use of store-and-forward routing, as opposed to cut-through routing. Cyclic queuing and forwarding operates by collecting data frames with reserved bandwidth within a cycle as "prioritized" and re-transmitting them during the next cycle. This ensures that if a particular end-to-end transmission needs to be forwarded by two intervening switches, that the end-to-end delay will be limited to two cycle times, as shown in Figure 2.

For applications such as those where real time video is transmitted, cyclic queuing and forwarding provides a method for achieving zero congestion loss and highly predictable latency; i.e., low latency jitter. For video, this provides the advantage of minimizing the buffer memory size requirements for the receiving end point. For some applications, it’s possible to improve latency performance further by using frame preemption (section 2.5) in conjunction with cyclic queuing and forwarding.

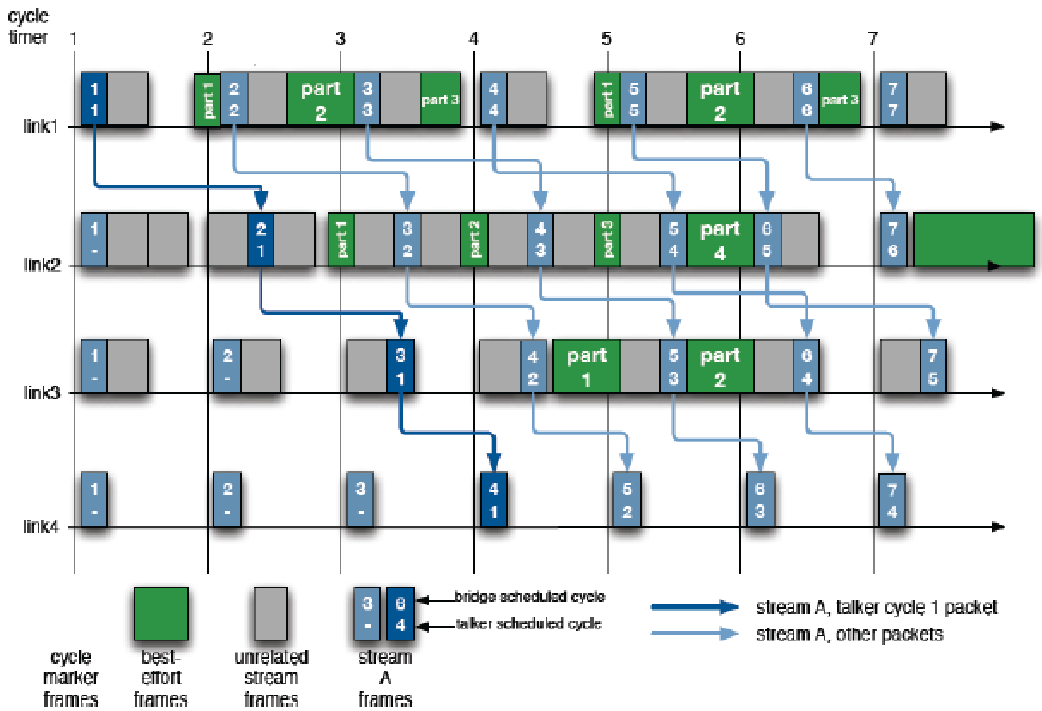

Figure 2. Cyclic Queuing and Forwarding

2.5 Frame Preemption

Certain applications such as control systems require extremely low latency jitter and therefore highly accurate arrival times. To support these applications, IEEE standard 802.1Qbu, Frame Preemption, provides an additional mechanism for prioritization.

Using frame preemption, the frames for individual streams are designated as being either “express” (higher priority) or “preemptible” (lower priority). Figure 3 illustrates the effect of frame preemption. In the top diagram of Figure 3(a), an interfering preemptible frame (possibly a best effort traffic frame) infringes past the scheduled transmission start time for an express frame. This time is indicated by the thick black vertical line in the diagram.

Referring to the lower timing diagram in Figure 3(a), one way to prevent this is by scheduling such that the interfering preemptable frame doesn’t transmit during the indicated “guard band” time. This will allow the scheduled express frame to begin its transmission on-time.

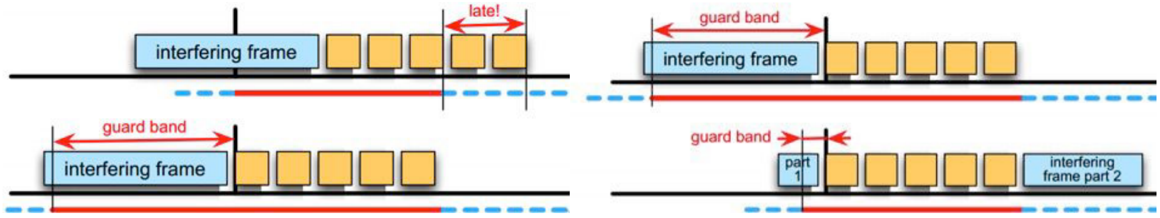

(a) (b)

Figure 3. Frame preemption: (a) Critical traffic delayed by interfering lower priority frame; (b) Critical traffic forwarded on-schedule by truncating interfering frame

The lower diagram in Figure 3(b) illustrates a second method for ensuring that the high priority express frame is transmitted on schedule. With this method, the transmission of the low-priority interfering frame is preempted by truncating its transmission a portion of the way through the frame. This allows the higher priority frame(s) to then be transmitted in its entirety on schedule. Following transmission of the higher priority express frame(s), the remainder of the interfering frame is then transmitted.

The advantage of frame pre-emption relative to the use of guard bands is that pre-emption makes optimal use of available bandwidth on the outgoing link of the switch or end point.

2.6 Time Synchronization

TSN's scheduling algorithms require a common knowledge of time by all end points and switch ports across a network. This requires a highly accurate time synchronization method. Methods for doing this are defined by IEEE standard 802.1AS, Timing and Synchronization. This standard is a profile for Ethernet based on an earlier standard, IEEE-1588. IEEE-1588 is the Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems, also referred to as Precision Time Protocol (PTP).

IEEE-1588 is a mature standard that's been widely used since 2008. The 802.1AS Timing and Synchronization standard is a defined profile of IEEE-1588.

Referring to Figure 4, IEEE 1588 and 802.1AS define a master/slave protocol used to synchronize all clocks on a network. This involves the use of one network node that includes a "Grandmaster" clock source, with all other nodes containing "boundary clocks" or clock slaves.

IEEE 1588 and 802.1AS may be used to provide time resolution and accuracy to the nanosecond range. For a network transmitting high data rates (such as 10 to

100 Gb/s data), this level of time accuracy will be necessary in order to enable accurate message scheduling and prevent traffic congestion.

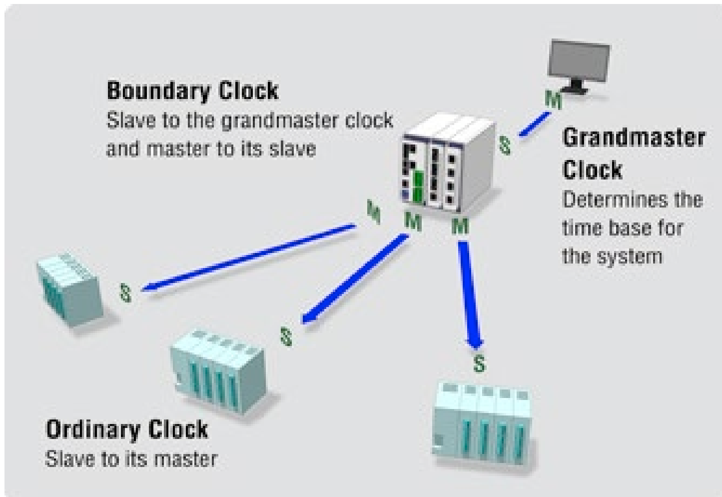

Figure 4. Time Synchronization

Some features of IEEE 1588 and 802.1AS time synchronization include:

- Common services for measuring link delays and rate ratios (clock frequency differences between nodes).
- A one-step clock, requiring only a single event message for providing time information.
- A mechanism for selecting the best master clock to be the grandmaster.
- Methods defining use and switchover for a back-up grandmaster.

Tightly synchronized time is necessary for determining the cycle times for cyclic queuing and forwarding. In addition to TSN traffic scheduling, the synchronized time is also usable for time stamping of application data such as synchronizing the timing of multiple incoming streams and marking the times that particular events occur.

2.7 Redundancy

For critical messages, it can be desirable to establish redundant paths through a network in order to avoid packet loss. For TSN, redundancy is defined by IEEE standard 802.1CB, Frame Replication and Elimination for Reliability.

Figure 5 illustrates an example of the use of redundancy by providing two paths between source node B4 and destination node B2. There are multiple methods for sending redundant messages:

1. Send all messages over path F. If a message fails over path F, then retry the message over path F'. This method has the advantage of reducing power consumption.

2. Send all messages over both path F and path F'. When node B2 receives a valid message over either of the two paths, it will then ignore a valid message received over the alternate path. If an invalid message is received, a subsequent valid message will then be received and accepted. Although this increases power consumption, it eliminates the need to retry messages over the alternate path. When operating over long distances, this approach could be advantageous.

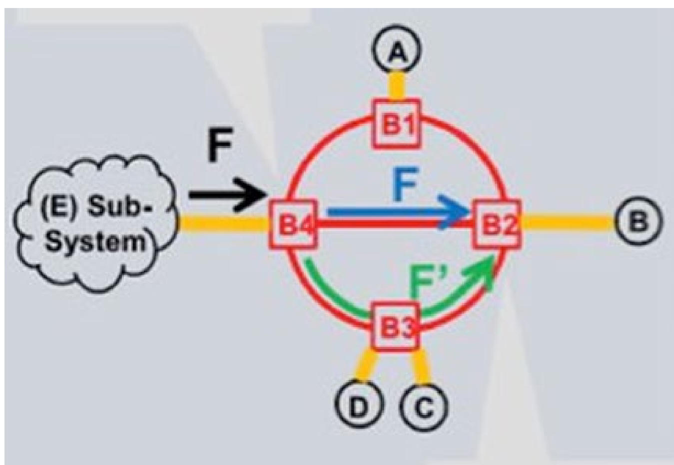

Figure 5. Redundant Frame Transmission

2.8 Filtering and Policing

IEEE standard 802.1Qci, Per-Stream Filtering and Policing, defines methods for protecting switches and end nodes against hardware errors, software bugs or hostile devices. Figure 6 illustrates the forwarding process functions defined by this standard.

There's a separate gate associated with each data stream. These gates are normally open, however they may be closed following detection of various types of violations. Each stream gate may be programmed to block traffic based on pre-defined criteria. In addition, each stream gate may include one or more counters. The counters are used to provide bandwidth information by providing “flow meters” for the respective streams.

The 802.1Qci standard enables filtering to be performed on a per-stream basis and also includes an option for switch policing. If used, the policing feature will result in a switch discarding frames received from an end node that’s sending packets and/or data at a higher rate than it’s allowed to. This is typically used for rate constrained traffic and provides a fail-safe mechanism to protect switches and/or downstream end nodes from “babbling” end nodes.

Other possible filters may be implemented based on priorities and traffic classes, and may be used to provide security. The latter could entail checking packets based on addresses, Stream IDs, labels, message sizes and payload content.

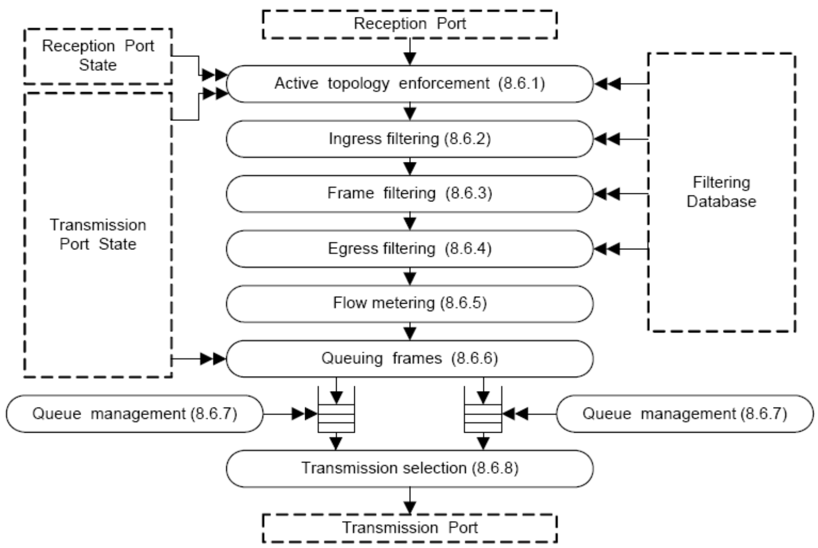

Figure 6. IEEE 802.1Qci Forwarding Process Functions

2.9 TSN Network Configuration

TSN includes four standards for defining how networks are configured. These are:

- P802.1Qcc, Stream Reservation Protocol (SRP) Enhancements and Performance Improvements
- 802.1Qca, IS-IS Path Control and Reservation
- 802.1Qat, Stream Reservation Protocol
- P802.1Qcc, Stream Reservation Protocol (SRP) Enhancements

These standards define methods for communicating and computing various network parameters. The latter include MAC addresses, time-aware scheduler configuration and parameters, class of service priorities for real time data streams, time slots for real time data, use of preemption, and redundancy paths.

There are three models for implementing TSN configuration:

1. A **fully centralized model**, referencing the top diagram in Figure 7. Under this model, talker and listener end points communicate their stream requirements to a logical centralized network configuration (CNC). This is done using a protocol called OPC UA (Open Platform Communications Unified Architecture) and possibly using a modeling language called IETF YANG/NETCONF.

   After the CNC collects this information, it calculates time slots, assigns resource reservations, and then configures all network participants. This includes configuring network switches using management protocols such as SNMP.

2. A **decentralized model**. Referencing the middle diagram in Figure 7, this model is based on the Stream Reservation Protocol (SRP) used with the earlier AVB standard. Talkers announce their streams to transmit, and listeners can subscribe to the individual streams. The network then reserves the necessary resources and time slots. The determination about whether the requirements can be satisfied is computed in a decentralized way within the network. This method does not require a central management entity.

3. A **partially centralized hybrid approach**. Referencing the bottom diagram in Figure 7, this approach involves a combination of the centralized and de-centralized methods, and uses a centralized network manager. With this method, end devices communicate with the nearest switch which forwards to the central management system using a standardized protocol.

The IEEE 802.1Qca standard is for IS-IS (Intermediate System to Intermediate System) Path Control & Reservation (PCR). It provides IS-IS control beyond shortest path trees (SPTs) by also including non-shortest path capabilities. IS-IS provides basic functions, including topology discovery and default paths, and involves one or more controllers controlling Explicit Trees.

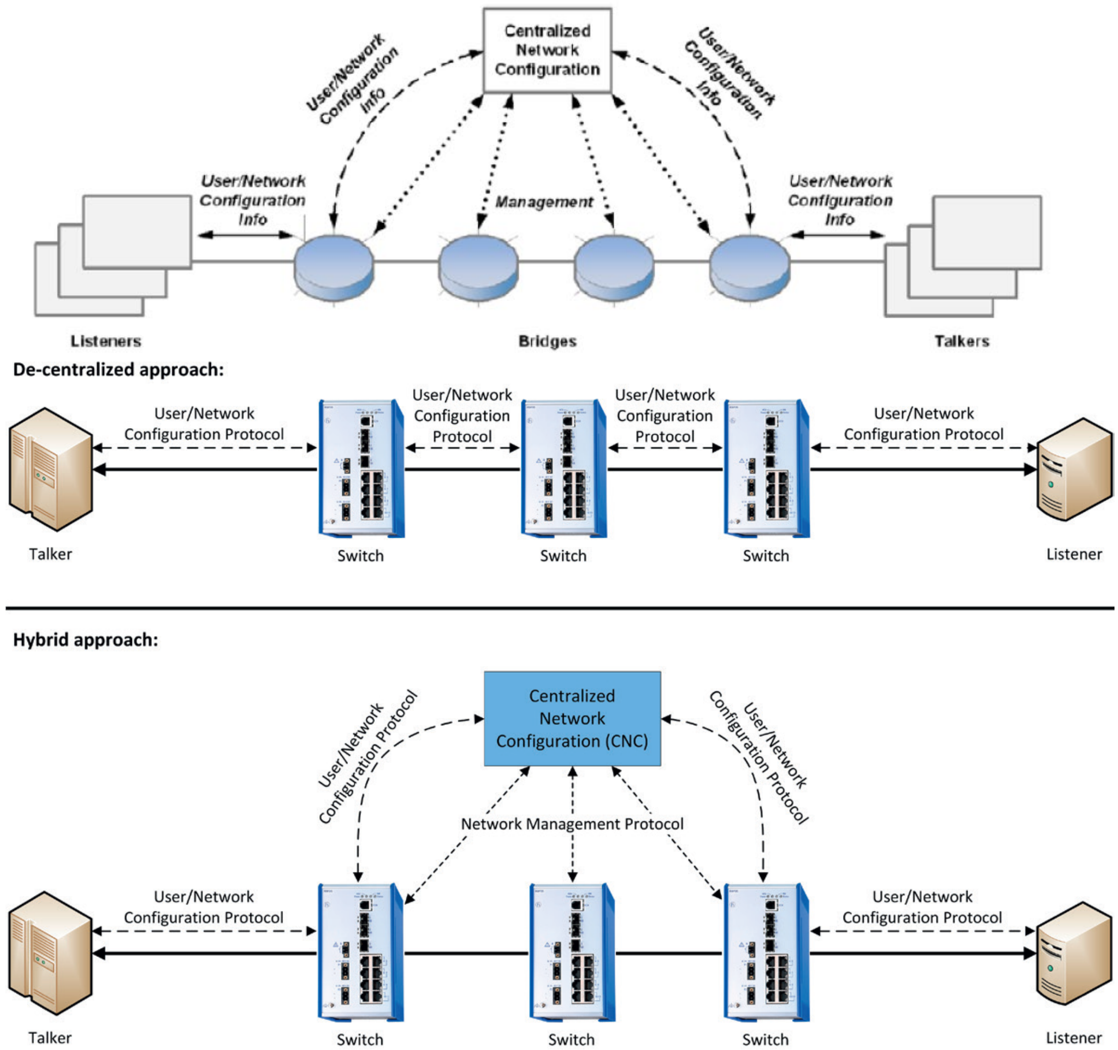

Figure 7. Centralized, Decentralized and Hybrid Network Configuration Models

3 Summary/Conclusion

TSN's implementation of time synchronization, deterministic traffic scheduling, redundancy, and policing provide a robust solution for extending the capabilities of Ethernet to address a number of core requirements within aerospace systems. TSN contains an ideal framework for the development of a certifiable implementation of deterministic Ethernet based on a profiled subset of IEEE 802.1Q. The envisioned aerospace profile would be derived from open standards and would provide the aerospace industry with a solution that leverages best practices from a larger commercial networking market.